Think Cloud Compliance

An Introduction to Cloud Computing for Legal and Compliance Professionals

What is cloud computing? Cloud computing—also known as simply "the cloud"—is a well-established concept in IT departments based on terms and definitions developed by the US government-based National Institute of Standards and Technology (NIST) and the non-governmental International Organization for Standardization (ISO). With seven service models and four deployment models, it has become an increasingly prevalent means for organizations to reduce information technology (IT) costs and complexity while increasing flexibility to rapidly deploy new IT workloads. Yet for non-technical professionals such as business managers and in-house legal staff, it may not be clear exactly what cloud computing is or how it differs from more traditional computing models such as self-managed data centers or outsourced IT services.

Any company—from a small shop to a worldwide enterprise corporation—can cut costs and alleviate the complexity of onsite IT management by shifting data and applications to a cloud service provider (CSP). While transferring data to a cloud vendor may ease the burden of data storage and application management for the small business owner or the enterprise IT department, it raises questions about data security, privacy, service reliability and legal and regulatory compliance. When considering the cloud, both the small business owner and the enterprise legal compliance officer must ask themselves these questions: Can a cloud provider keep data more secure than a company's own onsite system? Where is cloud data stored? Who has access to it? How reliable are cloud providers? Can a cloud provider maintain compliance with standards and regulations required for multi-national organizations with businesses in different verticals such as healthcare or education?

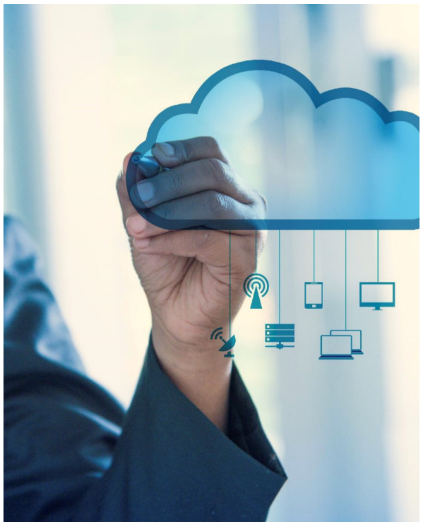

The goal of this paper is to provide some basic concepts, definitions and examples related to cloud computing for non-technical audiences as well as to address the questions that are raised by the shift to cloud computing.

Background

The cloud is no longer a future concept or a special computing model only appropriate for tech savvy companies. In fact, cloud computing has become a standard way of doing business, perfectly suited for all types and sizes of companies that are overwhelmed by their current computer systems and looking for ways to off-load IT complexity. According to Gartner, 89% of companies were using cloud computing in some form by the end of 2016. Many companies may be using the cloud today without realizing it. For example, if a company subscribes to online services such as Dropbox, Box, Gmail, Office 365 or SalesForce then they are already "in the cloud." Whether a business is contemplating a move to the cloud, or has perhaps begun the migration unknowingly, it is important that the company maintain the same level of security, privacy and regulatory compliance of its data that it had when the data was controlled onsite.

But how can businesses be certain that those levels are maintained once data leaves their premises? Because the same regulations and standards designed to protect the electronic transmission of personal data when data were only being transferred directly between providers and consumers apply to intermediary cloud providers as well.

For example, any business in the US that handles electronic health care records is bound by HIPAA regulations regardless of whether it stores data on-premises or with a cloud provider. When the "simplified" HIPAA administration document is 115 pages long, how can a company that primarily provides health care possibly have the expertise to understand the full HIPAA administrative document? Yet, that is the expectation for any company that handles electronic health care records. Now imagine an international company handling electronic health records, student data and criminal justice information. Each of those industries in each country is bound by different sets of regulations and standards. Such a company would either need a huge internal legal and compliance department or trust its data to a cloud computing provider that would have dedicated resources to meet both industry-specific data handling regulations as well as federal and international data handling laws.

According to a recent Forrester survey[1], three years ago the main reason businesses moved to the cloud was to save money and reduce complexity. Now, the most important reason is a desire to improve security and the ability to meet compliance regulations. Both drivers are really about reducing an organization's risk profile by transferring some of the risk to an enterprise cloud vendor who is better equipped to deal with security and compliance.

What is the cloud?

Whatever motivation a company has for moving to the cloud, let's take a step back and review what cloud computing is, as well as computing methods that may look like cloud computing but are not.

There are two primary standards organizations that have developed terms and definitions for cloud computing—the US government-based National Institute of Standards and Technology (NIST)[2] and the non-governmental International Organization for Standardization (ISO). The formal cloud computing definition from ISO/IEC 17788:2014[3] is as follows:

> *“Cloud computing: Paradigm for enabling network access to a scalable and elastic pool of shareable physical or virtual resources (examples of resources include servers, operating systems, networks, software, applications, and storage equipment) with self-service provisioning and administration on-demand.”*

NIST
National Institute of
Standards and Technology
U.S. Department of Commerce

ISO IEC
ISO/IEC JTC 1
Information technology

Essential Characteristics

These organizations have gone further by developing essential characteristics that more precisely define cloud computing. These characteristics also represent some of the key benefits of cloud computing. The six essential characteristics as stated in the ISO standard are:

- ✓ ***Broad network access***: *A feature where the physical and virtual resources are available over a network and accessed through standard mechanisms that promote use by heterogeneous client platforms. The focus of this key characteristic is that cloud computing offers an increased level of convenience in that users can access physical and virtual resources from wherever they need to work, as long as it is network accessible, using a wide variety of clients including devices such as mobile phones, tablets, laptops, and workstations.*
- ✓ ***Measured service***: *A feature where the metered delivery of cloud services is such that usage can be monitored, controlled, reported, and billed. This is an important feature needed to optimize and validate the delivered cloud service. The focus of this key characteristic is that customers only pay for the resources that they use. From the customers' perspective, cloud computing offers the users value by enabling a switch from a low efficiency and asset utilization business model to a high efficiency one.*
- ✓ ***Multi-tenancy***: *A feature where physical or virtual resources are allocated in such a way that multiple tenants and their computations and data are isolated from and inaccessible to one another. Typically, and within the context of multi-tenancy, the group of cloud service users that form a tenant will all belong to the same cloud service customer organization. There might be cases where the group of cloud service users involves users from multiple different cloud service customers, particularly in the case of public cloud and community cloud deployments. However, a given cloud service customer organization might have many different tenancies with a single cloud service provider representing different groups within the organization.*
- ✓ ***On-demand self-service***: *A feature where a cloud service customer can provision computing capabilities, as needed, automatically or with minimal interaction with the cloud service provider. The focus of this key characteristic is that cloud computing offers users a relative reduction in costs, time, and effort needed to take an action, since it grants users the ability to do what they need, when they need it, without requiring additional human user interactions or overhead.*
- ✓ ***Rapid elasticity and scalability***: *A feature where physical or virtual resources can be rapidly and elastically adjusted, in some cases automatically, to quickly increase or decrease resources. For the cloud service customer, the physical or virtual resources available for provisioning often appear to be unlimited and can be purchased automatically in any quantity, at any time, subject to constraints of service agreements. Therefore, the focus of this key characteristic is that cloud computing means that customers no longer need to worry about limited resources and might not need to worry about capacity planning.*
- ✓ ***Resource pooling***: *A feature where a cloud service provider's physical or virtual resources can be aggregated in order to serve one or more cloud service customers. The focus of this key characteristic is that cloud service providers can support multi-tenancy while at the same time using abstraction to mask the complexity of the process from the customer. From the customers' perspective, all they know is that the service works, while they generally have no control or knowledge over how the resources are being provided or where the resources are located. This offloads some of the customers' original workloads, such as maintenance requirements, to the provider. Even with this level of abstraction, it should be pointed out that users might still be able to specify location at a higher level of abstraction (e.g., country, state, or data center).*

Along with the essential characteristics, the NIST and the ISO have defined several cloud computing service models and sub-models as well as four deployment models. The first three service models are the primary models on which the other models are based.

Service Models

The following cloud service categories provide capabilities to the cloud service customer (in this case the NIST definitions are in italics):

- **Infrastructure as a Service (IaaS)**: *The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over the operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).*
  The most common example of IaaS is running virtual machines in a remote cloud environment. While the customer has control over the rapid provisioning and management of the virtual machine hosted operating environment, the CSP manages the infrastructure that the virtual machine runs in as well as supporting services such as networking and storage. The CSP also ensures there is a large enough pool of resources to accommodate the new infrastructure requirements.

- **Platform as a Service (PaaS)**: *The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.*
  In this case, the customer or an independent software vendor (ISV) creates a computer program that executes in the cloud using a set of application programming interfaces (APIs) specifically tuned to executing in a cloud environment. Instead of the program running on a local computer, the program runs in the cloud and provides services to users and client processes that interface with the program.

- **Software as a Service (SaaS)**: *The capability provided to the consumer to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a Web browser (e.g., Web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.*
  The most common examples of a SaaS solution is cloud-based email as a service, such as Microsoft Exchange Online (which is part of Office 365) or Google Gmail. In both cases, the user can interface with the email service using a Web browser connected to the cloud service or they can configure a mail application on a client device, such as Outlook, to interface with the service and download and send email using the application's interface.

In addition, there are numerous sub-models that have been offered by CSPs and we can expect new models to appear in the future. The additional sub-models that the ISO defines are as follows:

- **Communications as a Service (CaaS)**: This model includes real-time communications, interaction and collaboration services and can be delivered as PaaS or SaaS. Microsoft's Skype for Business is an example of this type of service.

- **Compute as a Service (CompaaS)**: The ISO defines this as the "*provision(ing) and use of processing resources needed to deploy and run software.*" This service is delivered as IaaS and is actually the original cloud service model. In the early days of Amazon's EC2 offering, customers could buy raw compute power or server capacity and be billed using a consumption-based model.

- **Data Storage as a Service (DSaaS)**: The ISO defines this as the "*provision(ing) and use of data storage and related capabilities*" and it can be delivered using IaaS, PaaS or SaaS. The most common examples of DSaaS would be virtual data storage offerings such as Dropbox, Box, Google Drive and Microsoft OneDrive. In each case, customers can connect to the cloud service and use the cloud to store files.

- **Network as a Service (NaaS)**: The ISO defines this as the delivery of "*transport connectivity and related network capabilities.*" This service can be delivered as IaaS, PaaS or SaaS and can address network enhancement, security and bandwidth challenges. A common example of NaaS would be a virtual private network (VPN) service offered by a CSP. The customer connects to the CSP's network infrastructure and the cloud offering then tunnels the customer's network traffic using the cloud-delivered VPN service.

Cloud deployment models

Cloud deployment models represent how cloud computing can be organized based on the control and sharing of physical or virtual resources.

The cloud deployment models, as defined by the ISO, include:

- **Public cloud**: *Cloud deployment model where cloud services are potentially available to any cloud service customer and resources are controlled by the cloud service provider. A public cloud may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud service provider. Actual availability for specific cloud service customers may be subject to jurisdictional regulations. Public clouds have very broad boundaries, where cloud service customer access to public cloud services has few, if any, restrictions.*
  Public cloud is the most common deployment model. Examples include Amazon AWS, Google G Suite and Microsoft Azure and Office 365.

- **Private cloud**: *Cloud deployment model where cloud services are used exclusively by a single cloud service customer and resources are controlled by that cloud service customer. A private cloud may be owned, managed, and operated by the organization itself or a third party and may exist on premises or off premises. The cloud service customer may also authorize access to other parties for its benefit. Private clouds seek to set a narrowly controlled boundary around the private cloud based on limiting the customers to a single organization.*
  Private clouds are less common but popular with very large organizations that want more control over the cloud infrastructure. For example, IBM offers private cloud services and Microsoft offers the packaged Azure Pack software solution to allow customers to front-end their own data center systems and services as cloud services within their own organizations. OpenStack is another tool that is popular for building private clouds.

- **Community cloud**: *Cloud deployment model where cloud services exclusively support and are shared by a specific collection of cloud service customers who have shared requirements and a relationship with one another, and where resources are controlled by at least one member of this collection. A community cloud may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises. Community clouds limit participation to a group of cloud service customers who have a shared set of concerns, in contrast to the openness of public clouds, while community clouds have broader participation than private clouds. These shared concerns include, but are not limited to, mission, information security requirements, policy, and compliance considerations.*
  Community clouds are most popular in government or law enforcement scenarios where there is a need to share data and resources across similar agencies and a requirement to isolate the environment from public access.

- **Hybrid cloud**: *Cloud deployment model using at least two different cloud deployment models. The deployments involved remain unique entities but are bound together by appropriate technology that enables interoperability, data portability and application portability. A hybrid cloud may be owned, managed, and operated by the organization itself or a third party and may exist on premises or off premises. Hybrid clouds represent situations where interactions between two different deployments may be needed but remained linked via appropriate technologies. As such the boundaries set by a hybrid cloud reflect its two base deployments.*
  Hybrid cloud is becoming more popular as organizations look to leverage existing data center resources wrapped as private cloud services set up to interact with public cloud services such as a SaaS email service or an IaaS virtual machine. Microsoft offers the Microsoft Azure Stack which allows customers to automate the integration of private and public cloud services while providing a single combined experience for users.

What’s not cloud computing

While the NIST and the ISO definitions encompass a variety of computing configurations, there are services that would only be considered “cloud computing” based on how they are delivered.

The following scenarios would not be considered true cloud services unless they were implemented with all six essential characteristics present:

- **Out-sourced IT or remote data centers**: Many IT vendors blur the definition of outsourced IT and cloud computing with some claiming outsourced IT is the same as a private cloud. Just because someone else manages your data center, doesn’t mean the services are hosted in the cloud. For example, true cloud services still need a self-service component, they need to be automatically scalable based on utilization and the customer should be charged based on cycles, bandwidth and storage used.
- **Virtual machine hosting**: While virtual machine hosting is a key workload in the cloud, simply hosting a virtual machine at a remote data center does not provide all the benefits that accrue from running a virtual machine in the cloud. For example, the customer should be able to create and retire virtual machines based on typical workload templates with a few simple commands and could instantly gain access to the newly-created environment. The customer should not have to worry about whether the host environment has the necessary resources to host a single virtual machine or a thousand virtual machines. The cloud provider should be able to automatically scale to meet the needs of the virtual machines.
- **Remote login or remote desktop**: Hosting physical desktops or servers at a remote site where users can log in to use the computers is just remote hosting. At the remote hosting site, someone needs to configure and manage those individual machines. Even if that process is automated, the ability to quickly scale up and scale down would be resource intensive. For this reason, virtually all cloud workloads run as virtual workloads or services which are easily spun up as customers need them.
- **Web-based applications or sites**: Many web-based services may appear to look very much like a cloud service but behind the scenes the site is running on a fixed number of machines and other resources. If demand for the service grows, a non-cloud service may not be able to scale to meet the demand and overall performance slows down for all users. A true cloud application or site would be able to take advantage of a large pool of resources and should be able to scale to meet any reasonable demand from users without any degradation of service.
- **Internet-based email**: Web-based email services were some of the first Software-as-a-Service cloud offerings from the big vendors such as Microsoft Office 365 and Google Gmail. While these offerings are true cloud services, other web-based email solutions may not feature the same self-service capabilities, nor can they be multi-tenant nor able to easily scale as more customers use the service. In addition, some non-cloud services are hosted in one data center in one location and therefore are not capable of failing over to another location or may have poor performance if accessed from afar.
- **Client-server computing or distributed computing**: Client-server computing was popular before the rise of cloud computing. For example, many retail operations would have basic client devices for service personnel to use when dealing with customers and these devices would interact with servers at a remote location and exchange data with those servers, such as record a sale. While it might appear that the servers were “in the cloud,” they were simply somewhere else and most of the essential characteristics—such as rapid elasticity, scalability and resource pooling were not available. Other characteristics such as self-service were also not available. For example, adding a new point of sale system typically required service personnel visits and extensive manual configuration.

Again, for a service to be classed as a true cloud computing service, all six of the essential characteristics need to be present.

Cloud deployment model examples

With the cloud definitions out of the way and some clarity on what constitutes a cloud or not, let's look at a couple of typical examples.

Multi-tenant public cloud

As this graphic example shows, there are three service models (SaaS, PaaS and IaaS) hosted by a cloud vendor and these solutions service three different companies and within those companies, multiple workers access the various services. Since multiple companies use the same services—but in an isolated way—this is considered a multi-tenant environment. For example, the SaaS solution might be an email service such as Gmail. Worker #1 from Company A is using the service at the same time as Worker #2 from Company B. They are both using the same service but they are totally isolated from each other and the user and company data are protected from being exposed to other customers using the same service. Other workers are building cloud apps using a PaaS solution and others are running virtual machines for hosting custom workloads. All these run isolated from each other as though each user had his or her own machine or dedicated application environment. If more users access the services, the cloud environment is engineered to automatically draw from its pool of resources and scale to meet the demands of the users.

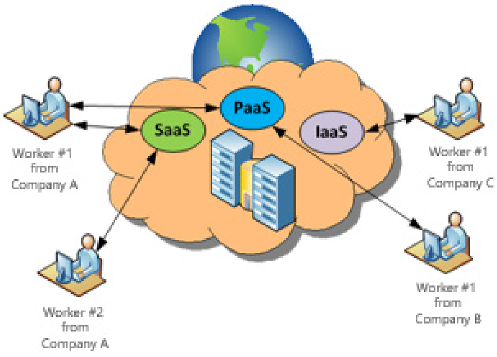

Figure 1 Multi-tenant public cloud example

Another key part of this example is the use of the term public cloud. This refers to the fact that these services are available for use by multiple customers and these customers typically gain access to the services using a secure connection via the public Internet infrastructure. Access could be simply a web-based application running in the user's browser as the interface into the service in the cloud (for example, a web-based email service such as Gmail) or it could be a rich client application on a PC or mobile device that front ends access to the backend services hosted in the cloud (for example, Outlook running on a PC accessing mail hosted in Office 365).

The cloud provider might also have multiple data centers around the country or around the globe in order to provide better performance or fault tolerance in the event of a disaster. To the users, there is no sense of this backend architecture as it is totally transparent to the user experience.

The key point in this example is multiple customers and users are sharing access to the same set of services and resources but each is totally protected and isolated from the other.

Single entity private cloud

One of the more confusing deployment models is the private cloud. In this example, there is a single company accessing services hosted in a cloud but this cloud is not shared with other companies. Not only is this cloud not shared, but the infrastructure is often built behind firewalls to isolate it from the public Internet. And in addition, the cloud data center is often hosted in a company facility or possibly a remote facility that the company or its agents control. These facilities are accessed using a VPN or similar secure network tunneling technology. The confusion arises as we consider if this is really a cloud scenario or is it simply a dedicated data center hosting applications for local and remote workers around the company. Again, our essential characteristics are the key to distinguishing this environment as a cloud or not. If the services are running on static servers, then this is probably not a cloud. If the services run as virtualized workloads that automatically scale to meet demand and can be easily created and retired by departments and users, then it

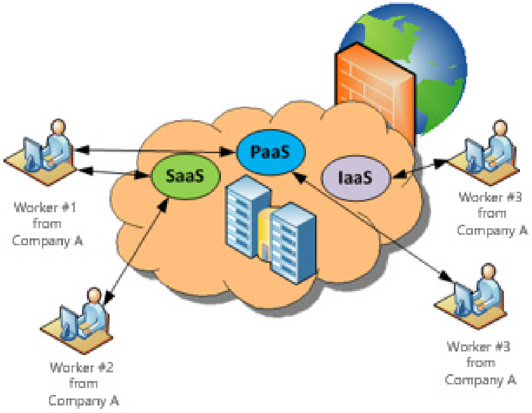

Figure 2 Single entity private cloud example

probably is a cloud environment. The fact that a single entity uses the environment does not impact whether this is a cloud or not. The elastic, highly-automated, resource-pooled architecture is the key. This architecture, while simple to use, is based on an extremely complex foundation and therefore is not something that most companies are in a position to build for themselves. For this reason, companies typically turn to off-the-shelf toolsets or packaged software offerings to build their private cloud environments. For example, Microsoft packages up the core software components that make up its Azure cloud services and makes it available as the Windows Azure Pack, which runs on top of Microsoft Server, System Center, SQL Server and more. To the users, this makes their own data center software look like cloud services and extends capabilities by offering self-service, scaling, virtual machine hosting and more.

Private clouds appear to be more popular in highly regulated industries where data protection is paramount or in other applications that deal with highly sensitive data. Some customers feel they are better protected if they control the entire environment and the physical security of their data centers. But doing security right can be very expensive and may defeat the cost benefits of cloud computing. Ultimately, large cloud providers such as Microsoft, Amazon and IBM can offer levels of security that far exceed what individual organizations can provide for themselves. Since the large cloud providers are protecting thousands of customers, their collective knowledge on how to deal with threats means they have a better chance to stay ahead of the bad guys. As a result, secure enterprise public clouds are likely to be more popular as customers become more comfortable with cloud computing.

On-premises vs. Cloud Computing

Even with these examples of cloud computing, it is worth reviewing the key differences between traditional "on-premises" computing and cloud computing. The term "on-premises" is a bit of a misnomer since a corporate data center might be in a different locality and users access the data center services using remote sessions, web-based sessions or other client-server technologies such as file sharing over the Internet or a VPN connection. Because some of these "on-premises" data centers are remote, this type of computing is often confused with cloud computing which also accesses remote services. So how do they differ?

On-premises computing

On-premises computing involves one or more data centers that are owned and/or controlled by the organization that they serve. The servers and equipment in the data centers are also owned or leased by the organization, all of which operates under a capital expenditure model. That is, the organizations purchase the assets with large upfront outlays of funds. In some cases, leasing is involved but the equipment is leased regardless of how much of the equipment resources are used. Organizations typically operate the data center or hire a third party to run it for them and they fully control the software and data in the data center. This includes installing operating systems, software and updates and ensuring the data center buildings and property and software are all properly secured. A firewall is typically used to prevent unauthorized access to the data center and security software is used to allow authorized remote access. In other words, operating your own data center can be very expensive because of large asset purchases, staffing for operations and building costs.

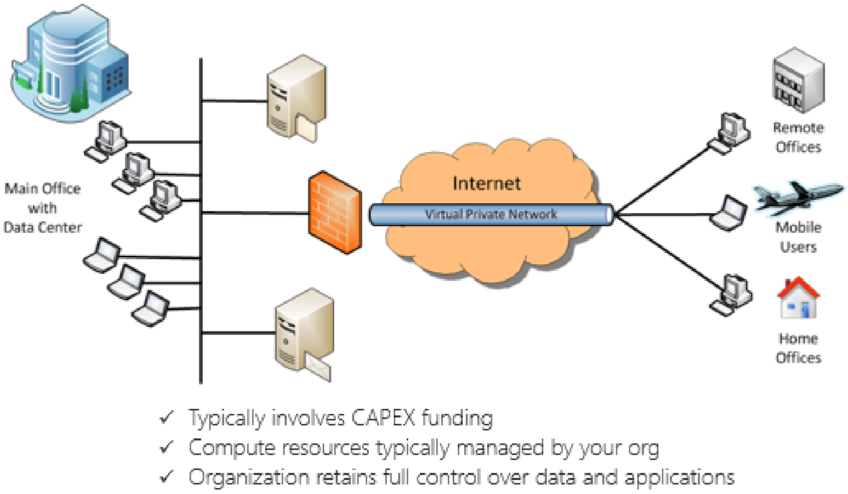

✓ Typically involves CAPEX funding
✓ Compute resources typically managed by your org
✓ Organization retains full control over data and applications

Figure 3 A representation of on-premises computing

Local users gain access to the data center using the onsite network. Remote users access the data center services using secure communications over the Internet or in some cases using dedicated long-haul network connections.

Cloud computing

In cloud computing, unlike on-premises computing, data center operations are all handled by the cloud service provider. Instead of high capital outlays and the need for data center staff, the organization operates under an operational expense model where departments and users are charged for the services they use, in a similar way to how an organization pays for electricity based on consumption. The cloud provider owns the software and makes it available as a pay-per-use service. These services can include access to virtual machines, software services such as email hosting and platform services that allow organizations to write applications designed specifically for the cloud. Even though the cloud provider controls the infrastructure, the users still own the data. Techniques such as encryption and tenant isolation keep data safe from other cloud users and even the cloud provider. This is an important point for customers in highly regulated markets such as healthcare or financial services or sensitive data scenarios such as law enforcement. Using the cloud should not compromise data protection obligations or requirements. Services can be hosted across a transparent network of cloud data centers in different countries and when a customer connects to the service they are typically unaware of which data center they are connected to. This allows users to get optimal service regardless of where they are connecting from.

Note that the user experience for both on-premises and cloud computing can be exactly the same. The main differences are with the backend infrastructure—both operationally and the funding model.

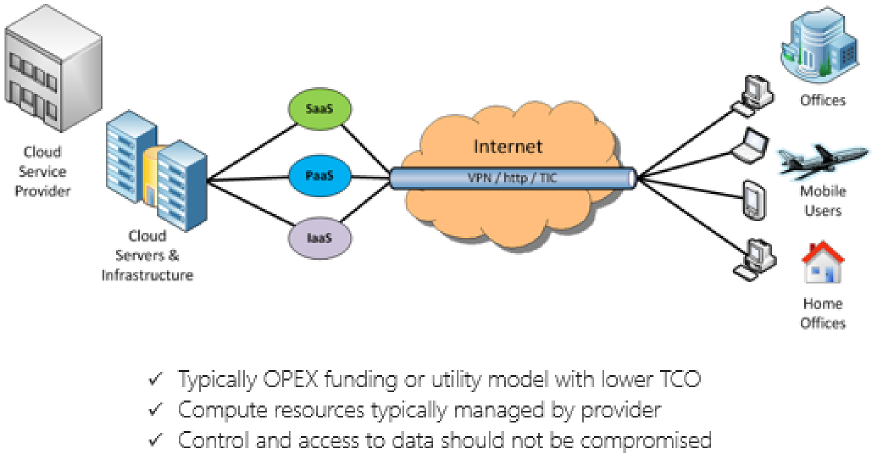

✓ Typically OPEX funding or utility model with lower TCO
✓ Compute resources typically managed by provider
✓ Control and access to data should not be compromised

Figure 4 A representation of cloud computing

Example: On-premises email vs. cloud-based email

To illustrate what this means, let's look at backend email services, both using an on-premises model with corporate-owned servers and using a cloud Software-as-a-Service model. In both cases, the user experience is again exactly the same. The user is running a mail client application, such as Microsoft Outlook, on his or her corporate PC as well as an email app on his or her personal mobile device. The user has no idea whether the back-end service is running on a corporate server or in the cloud.

With the email services running on a server in an organization's data center, the organization is responsible for acquiring the server and server software and hosting it in a secure location as well as managing the operation and updates for the server environment. This server might also be running other corporate workloads such as database applications. With multiple users and multiple workloads, the user experience can be variable, meaning high user load results in worse performance. The organization is also responsible for spam control, malware prevention and compliance with regulations that impact its business data. Overall, running this type of operation can be expensive, resource intensive and can provide an uneven user experience.

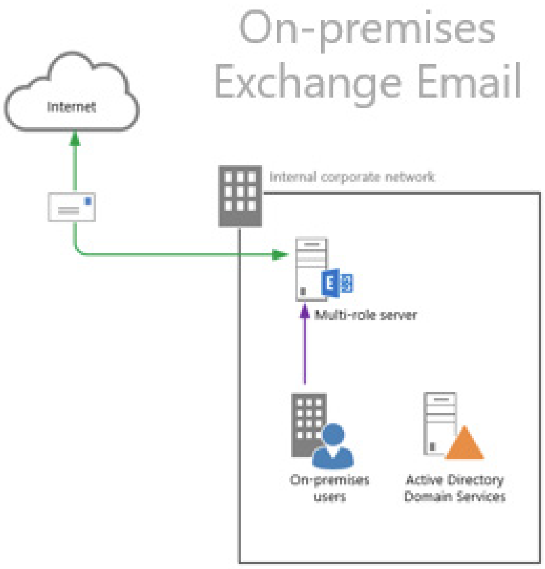

Figure 5 Example of on-premises email service

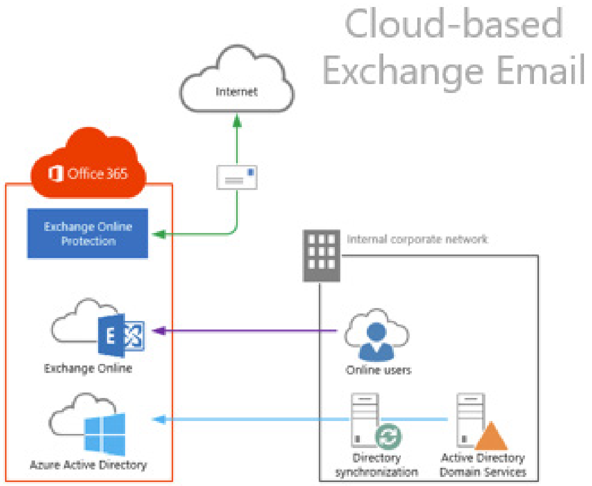
Cloud-based
Exchange Email
Internet
Office 365
Exchange Online Protection
Internal corporate network
Exchange Online
Online users
Azure Active Directory
Directory synchronization
Active Directory Domain Services

Figure 6 Example of a cloud-based email service

In a cloud scenario, the user email client connects to a cloud service instead of an email service running on a server in the corporate data center. If we look at Office 365 with Exchange Online as an example, the customer walks through some simple steps to configure email for his or her organization and the mailboxes for each user are auto-deployed. Authenticated access to the mailboxes can be through a cloud-based or on-premises identity management system. In this case, Microsoft provides access to the back-end resources in the form of cloud services. It handles the management, patching and updating of the systems and ensures the environment is secured—with the help of teams of hundreds of engineers, security experts and legal staff—and data processing is managed in compliance with applicable regulations and laws. Since Microsoft has a large pool of resources to draw from, the system automatically scales up or down to handle workloads with consistent performance. Microsoft also makes the services available transparently through local network entry points around the world so that the user experience is again as consistent as possible. The customer is only billed for the services used by each user and is not responsible for the purchase or maintenance of the back-end equipment and services.

For most organizations—banks, insurance companies, airlines, hospitals and governments—IT is not their primary focus. Cloud computing allows them to focus on their core business and lets the cloud provider focus on delivering the service in a secure and compliant way.

Where is the cloud?

Modern cloud providers use a modular system of servers, data storage and networking components that can be dropped into a data center anywhere in the world. These compute modules are connected to electricity for power and water for cooling and then they can be provisioned automatically to join the cloud.

The major cloud providers typically have data centers around the world to serve local users and to provide some level of redundancy in the event of disruptions in other data centers. In some cases, these data centers may only be available for a particular set of users based on national data residency requirements or regulations, or to address the special needs of a community of customers, such as government users. In other cases, using local data center trustees is a way for a cloud provider to have local government data requests handled by the legal team of the local data center trustee.

Ultimately an enterprise-grade cloud vendor should provide transparent, seamless, secure and fast access to its cloud services from almost anywhere in the world while complying the regulatory needs of each region and its customers.

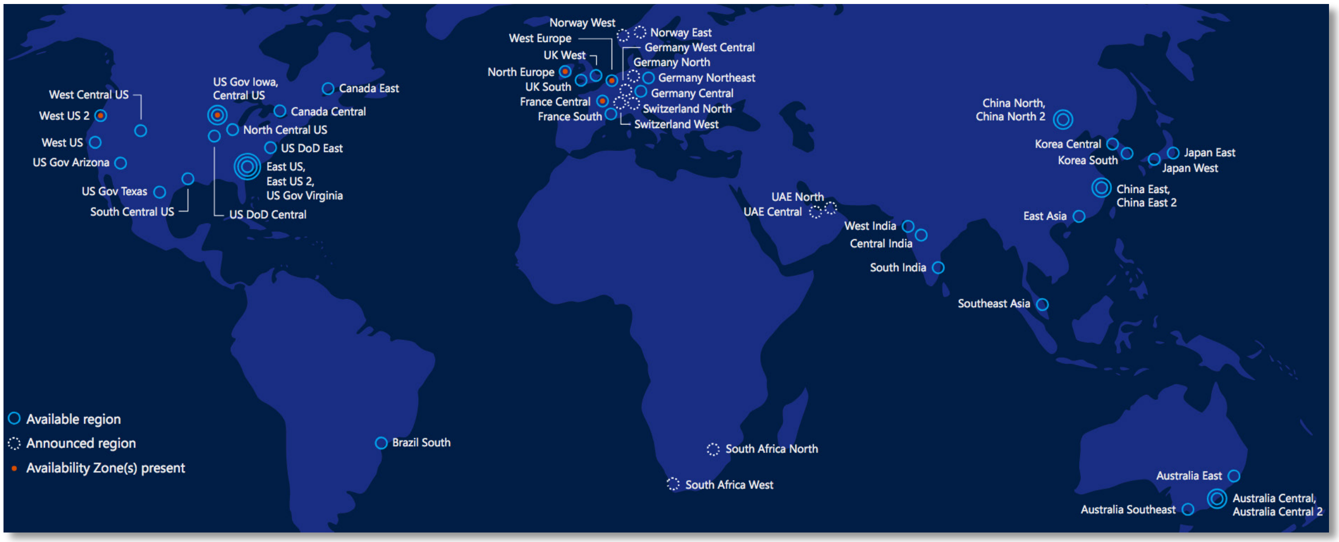

Figure 7 Microsoft's global network of cloud data centers

Compliance in the cloud

Managing compliance, especially for multi-national organizations, is a complex task that is difficult for an organization to navigate on its own. Compliance is even more of a challenge for regulated industries such as healthcare or financial services. Not only are there numerous standards and regulations, but these are constantly changing making it even more difficult for a business to keep abreast of international electronic data handling laws.

One of the biggest benefits of partnering with a well-established cloud vendor is that many of the regulatory and data protection obligations and requirements to comply with recognized standards can be handed off to, or at least shared with, the cloud vendor. Large cloud vendors address regulatory compliance needs every day and because these vendors are dealing with large volumes of customers and data that span virtually every industry and every country, their broad experience in standards and regulatory compliance is more complete than any one company can have on its own. When proper technical and legal frameworks are in place, cloud providers can handle their customers' most sensitive data and satisfy the needs of even the most stringent data protection regulations. Customers can know that their data are being handled properly if their cloud provider has met the criteria for these international standards.

For this reason, customers should evaluate the breadth and depth of standards and regulatory compliance for any cloud vendor being considered for sensitive data processing.

When we look at the ways in which cloud service providers handle compliance issues, we really have three broad areas of standards or regulations to consider.

Cross-industry international standards

The first group are standards from bodies such as the ISO or the Cloud Security Alliance which are applicable to virtually all cloud customers around the world. These have arisen out of demands from customers for a consistent approach to operations, security, privacy, risk management and governance. Certifications provide uniform methods for measuring a cloud provider's ability to meet these standards—often with the help of independent third-party auditors.

Figure 8 Examples of international standards

Standards spanning verticals and regions

The next set of standards are typically regulations that have been put in place for different industries. These industries include healthcare, manufacturing, education, financial services and government. In many cases, a company cannot offer online services to customers in these verticals unless it complies with the regulations created for that industry.

Finally, there are regulations and standards that based on regional or national needs or data protection laws. For example, US companies handling European citizen data are required to comply with the US-E.U. Privacy Shield and the European General Data Protection Regulation (GDPR).

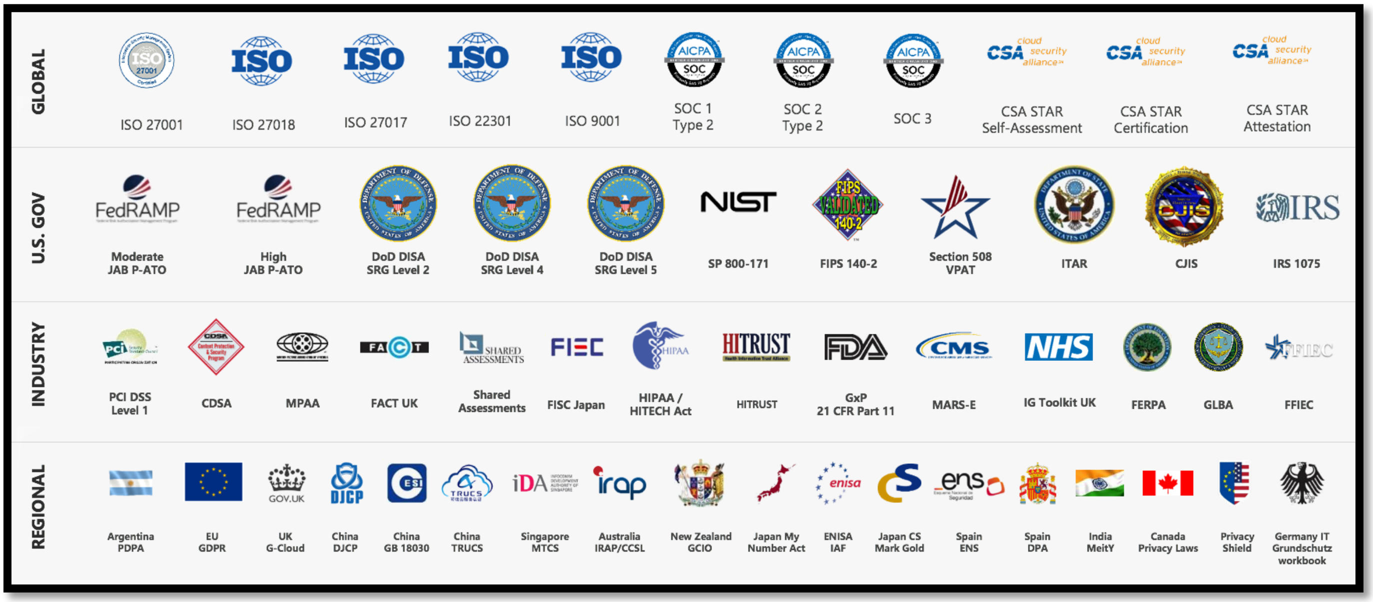

Figure 9 Examples of global, vertical and regional standards

What a Compliant Cloud Offers

Compliance is more complex than just meeting a check list of standards and regulations. Effective compliance requires a two-way partnership between the customer who owns the data along with the legal obligations for the handling of the data and the cloud vendor who acts as the data processor and must also handle the data in compliance with regulations. Having a partnership agreement that clearly defines roles and responsibilities is essential to achieve complete legal and regulatory compliance.

Standards also play a vital role in reassuring customers that the cloud vendor conforms with expected norms for security, breach prevention, privacy, operations and more. In many cases conformance with these standards is not just a pledge—a third party auditor is engaged to assess, monitor and validate the cloud operator.

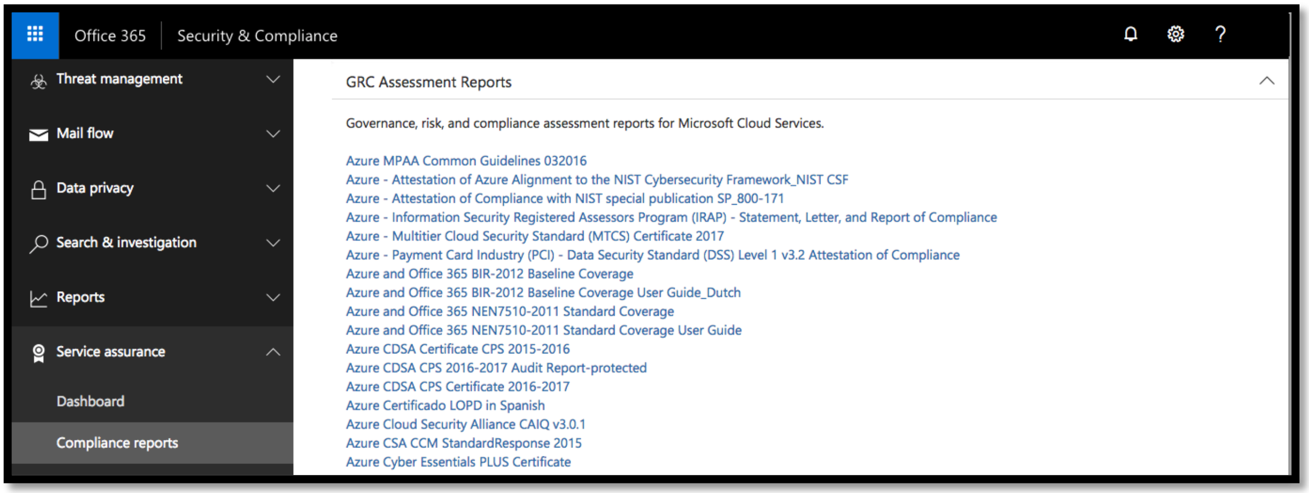

GRC Assessment Reports

Governance, risk, and compliance assessment reports for Microsoft Cloud Services.

Azure MPAA Common Guidelines 032016
Azure - Attestation of Azure Alignment to the NIST Cybersecurity Framework_NIST CSF
Azure - Attestation of Compliance with NIST special publication SP_800-171
Azure - Information Security Registered Assessors Program (IRAP) - Statement, Letter, and Report of Compliance
Azure - Multitier Cloud Security Standard (MTCS) Certificate 2017
Azure - Payment Card Industry (PCI) - Data Security Standard (DSS) Level 1 v3.2 Attestation of Compliance
Azure and Office 365 BIR-2012 Baseline Coverage
Azure and Office 365 BIR-2012 Baseline Coverage User Guide_Dutch
Azure and Office 365 NEN7510-2011 Standard Coverage
Azure and Office 365 NEN7510-2011 Standard Coverage User Guide
Azure CDSA Certificate CPS 2015-2016
Azure CDSA CPS 2016-2017 Audit Report-protected
Azure CDSA CPS Certificate 2016-2017
Azure Certificado LOPD in Spanish
Azure Cloud Security Alliance CAIQ v3.0.1
Azure CSA CCM StandardResponse 2015
Azure Cyber Essentials PLUS Certificate

Figure 10 Microsoft cloud customer portal for compliance reports

Cloud customers should have access to these audited reports, certifications and attestations. The legal staff in the company can use these documents to prove to their management that the cloud vendor they are dealing with addresses the organization’s regulatory compliance requirements. This is also a proof of compliance that can be used with the organization’s customers.

All of this should ultimately translate into tangible benefits for the cloud customers, including:

- More protection against data leaks and breaches
- Less risk of regulatory or legal sanctions
- Lower costs for achieving compliance
- Assurance of adherence to international privacy and security standards
- Respect for the rules of highly regulated industries
- Decreased overall risk for data and business
- Predictability for handling legal requests from law enforcement agencies

Given the ever-changing regulatory landscape, companies with businesses to run should have complete confidence trusting the security and privacy of their applications and data to a cloud provider with a proven commitment to compliance regulations and standards.

Trusting your cloud provider: Security and Compliance tools

For many organizations, the security and privacy of their data is their primary concern, with regulatory compliance simply a means to that end. How do you know if your cloud service is secure and private? Compliance with established security, privacy and data protection standards and regulations goes a long way towards addressing these concerns. Experienced enterprise cloud providers offer an easy way for legal and compliance professionals inside organizations to gain access to tools and information to help the customer’s staff operate in a consistent and compliant way. This includes access to governance, risk and compliance summaries and reports for the services in use as well as tools for managing employee security access settings, device management policies, data retention policies, data loss prevention and eDiscovery.

The cloud provider offers the tools, but the cloud customer must use them in order to operate a compliant cloud environment.

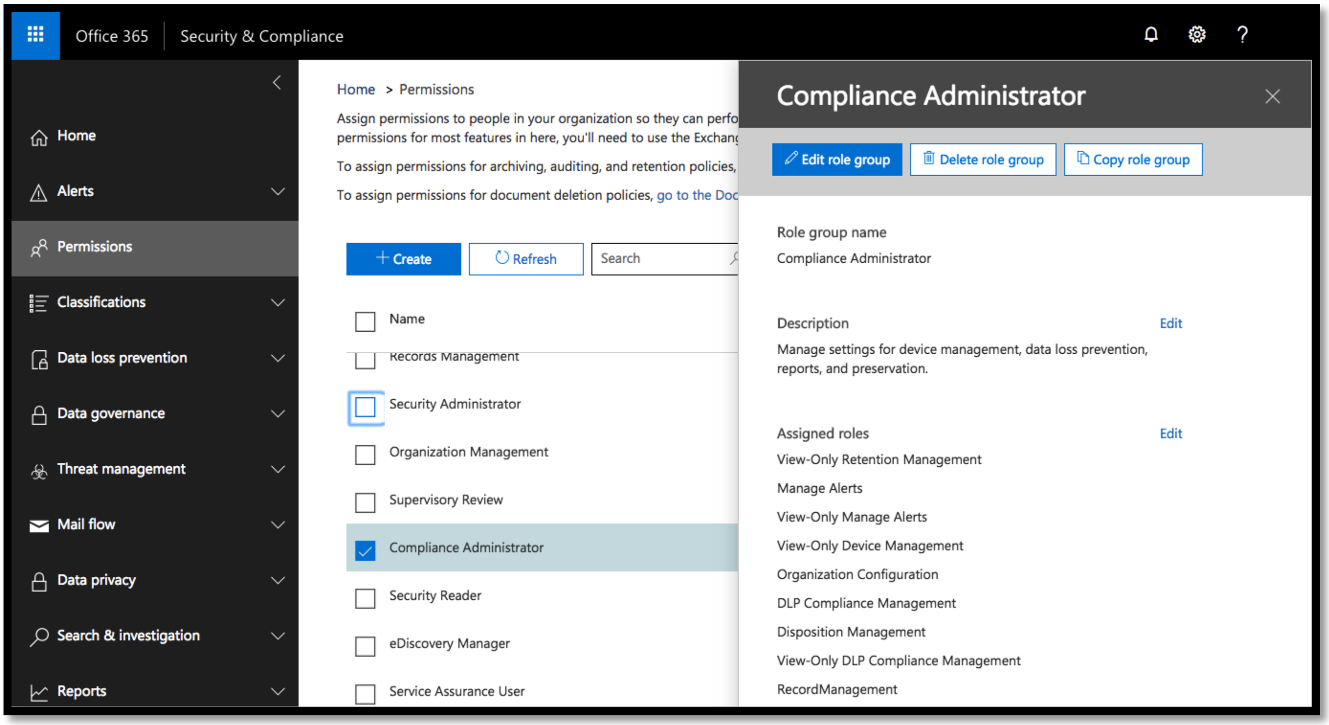

Home > Permissions

Assign permissions to people in your organization so they can perfo
permissions for most features in here, you'll need to use the Exchang

To assign permissions for archiving, auditing, and retention policies,

To assign permissions for document deletion policies, go to the Doc

□ Name
□ Records Management
□ Security Administrator
□ Organization Management
□ Supervisory Review
☑ Compliance Administrator
□ Security Reader
□ eDiscovery Manager
□ Service Assurance User

Compliance Administrator

Role group name
Compliance Administrator

Description
Manage settings for device management, data loss prevention, reports, and preservation.

Assigned roles
View-Only Retention Management
Manage Alerts
View-Only Manage Alerts
View-Only Device Management
Organization Configuration
DLP Compliance Management
Disposition Management
View-Only DLP Compliance Management
RecordManagement

Figure 11 Microsoft Security & Compliance portal for cloud customers

You can trust the cloud

Cloud computing has come a long way in the last ten years. In the early days, customers were attracted by the lower costs and reduced complexity of cloud computing but few customers—especially customers with sensitive data—were willing to take the risk of moving critical applications and data into the cloud.

Today, over 80% of enterprise customers use at least one cloud-based service and adoption is rapidly growing in both percentage penetration and the number of cloud workloads in use. Despite this growth in acceptance of cloud in the enterprise, trust is still the most important factor for selecting a cloud provider and a perceived lack of trust has historically been a reason for not going to the cloud. In a recent IDC study[4], over a third of cloud compliance decision makers ranked trust as the a most important factor with cost, flexibility and functionality less important. Trust factors are made up of security, privacy, compliance, reliability and transparency. Given concerns over data breaches, hacking, malware and other exploits, security is by far the most important trust factor for many organizations.

With security and trust being so important, how do potential cloud customers know for sure that they are dealing with a cloud provider and cloud technologies that can be trusted? Trust can be achieved by choosing a cloud provider who offers verified compliance with international, regional and industry standards and regulations.

Notes

[1] See “Study: Cloud Service Agreements Omit Key Considerations,” *Forrester Research* at https://aka.ms/forrester.iso19086

[2] See “US Government Cloud Computing Technology Roadmap Volume I,” *NIST* at http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.500-293.pdf

[3] See “ISO/IEC 17788:2014,” *ISO* at http://www.iso.org/iso/catalogue_detail?csnumber=60544

[4] “IDC-Microsoft LCC Survey Sept 2016”: Survey of 504 Cloud Compliance Decision Makers US-Germany-UK with power to veto or delay cloud deals on compliance grounds. Enterprises 1,000 or more employees.